

Memo to: Oaktree Clients
From: Howard Marks
Re: Is It a Bubble?

Ours is a remarkable moment in world history. A transformative technology is ascending, and its supporters claim it will forever change the world. To build it requires companies to invest a sum of money unlike anything in living memory. News reports are filled with widespread fears that America's biggest corporations are propping up a bubble that will soon pop.

During my visits to clients in Asia and the Middle East last month, I was often asked about the possibility of a bubble surrounding artificial intelligence, and my discussions gave rise to this memo. I want to start off with my usual caveats: I'm not active in the stock market; I merely watch it as the best barometer of investor psychology. I'm also no techie, and I don't know any more about AI than most generalist investors. But I'll do my best.

One of the most interesting aspects of bubbles is their regularity, not in terms of timing, but rather the progression they follow. Something new and seemingly revolutionary appears and worms its way into people's minds. It captures their imagination, and the excitement is overwhelming. The early participants enjoy huge gains. Those who merely look on feel incredible envy and regret and – motivated by the fear of continuing to miss out – pile in. They do this without knowledge of what the future will bring or concern about whether the price they're paying can possibly be expected to produce a reasonable return with a tolerable amount of risk. The end result for investors is inevitably painful in the short to medium term, although it's possible to end up ahead after enough years have passed.

I've lived through several bubbles and read about others, and they've all hewed to this description. One might think the losses experienced when past bubbles popped would discourage the next one from forming. But that hasn't happened yet, and I'm sure it never will. **Memories are short, and prudence and natural risk aversion are no match for the dream of getting rich on the back of a revolutionary technology that “everyone knows” will change the world.**

I took the quote that opens this memo from Derek Thompson's November 4 newsletter entitled “AI Could Be the Railroad of the 21st Century. Brace Yourself,” about parallels between what's going on today in AI and the railroad boom of the 1860s. Its word-for-word applicability to both shows clearly what's meant by the phrase widely attributed to Mark Twain: “history rhymes.”

Understanding Bubbles

Before diving into the subject at hand – and having read a great deal about it in preparation – **I want to start with a point of clarification. Everyone asks, “Is there a bubble in AI?” I think there's ambiguity even in the question. I've concluded there are two different but interrelated bubble possibilities to think about: one in the behavior of companies within the industry, and the other in how investors are behaving with regard to the industry.** I have absolutely no ability to judge whether the AI companies' aggressive behavior is justified, so I'll try to stick primarily to the question of whether there's a bubble around AI in the financial world.



The main job of an investment analyst – especially in the so-called “value” school to which I subscribe – is to (a) study companies and other assets and assess the level of and outlook for their intrinsic value and (b) make investment decisions on the basis of that value. Most of the change the analyst encounters in the short to medium term surrounds the asset’s price and its relationship to underlying value. That relationship, in turn, is essentially the result of investor psychology.

Market bubbles aren’t caused directly by technological or financial developments. Rather, they result from the application of excessive optimism to those developments. As I wrote in my January memo [On Bubble Watch](#), bubbles are temporary manias in which developments in those areas become the subject of what former U.S. Federal Reserve Chairman Alan Greenspan called “irrational exuberance.”

Bubbles usually coalesce around new financial developments (e.g., the South Sea Company of the early 1700s or sub-prime residential mortgage-backed securities in 2005-06) or technological progress (optical fiber in the late 1990s and the internet in 1998-2000). **Newness plays a huge part in this. Because there’s no history to restrain the imagination, the future can appear limitless for the new thing. And futures that are perceived to be limitless can justify valuations that go well beyond past norms – leading to asset prices that aren’t justified on the basis of predictable earning power.**

The role of newness is well described in my favorite passage from a book that greatly influenced me, *A Short History of Financial Euphoria* by John Kenneth Galbraith. Galbraith wrote about what he called “the extreme brevity of the financial memory” and pointed out that in the financial markets, “past experience, to the extent that it is part of memory at all, is dismissed as the primitive refuge of those who do not have the insight to appreciate the incredible wonders of the present.” In other words, history can impose limits on awe regarding the present and imagination regarding the future. In the absence of history, on the other hand, all things seem possible.

The key thing to note here is that the new thing understandably inspires great enthusiasm, but bubbles are what happen when the enthusiasm reaches irrational proportions. Who can identify the boundary of rationality? Who can say when an optimistic market has become a bubble? It’s just a matter of judgment.

Something that occurred to me this past month is that two of my best “calls” came in 2000, when I cautioned about what was going on in the market for tech and internet stocks, and in 2005-07, when I cited the dearth of risk aversion and the resulting ease of doing crazy deals in the pre-Global Financial Crisis world.

- First, in neither case did I possess any expertise regarding the things that turned out to be the subjects of the bubbles: the internet and sub-prime mortgage-backed securities. All I did was render observations regarding the behavior taking place around me.
- And second, the value in my calls consisted mostly of describing the folly in that behavior, not in insisting that it had brought on a bubble.

Struggling with whether to apply the “bubble” label can bog you down and interfere with proper judgment; we can accomplish a great deal by merely assessing what’s going on around us and drawing inferences with regard to proper behavior.

What's Good About Bubbles?

Before going on to discuss AI and whether it's presently in a bubble, I want to spend a little time on a subject that may seem somewhat academic from the standpoint of investors: the upside of bubbles. You may find the attention I devote to this topic excessive, but I do so because I find it fascinating.

The November 5 *Stratechery* newsletter was entitled “The Benefits of Bubbles.” In it, Ben Thompson (no relation to Derek) cites a book titled *Boom: Bubbles and the End of Stagnation*. It was written by Byrne Hobart and Tobias Huber, who propose that there are two kinds of bubbles:

. . . “Inflection Bubbles” – the good kind of bubbles, as opposed to the much more damaging “Mean-reversion Bubbles” like the 2000’s subprime mortgage bubble.

I find this a useful dichotomy.

- The financial fads I’ve read about or witnessed – the South Sea Company, portfolio insurance, and sub-prime mortgage-backed securities – stirred the imagination based on the promise of returns without risk, but there was no expectation that they would represent overall progress for mankind. There was, for example, no thought that housing would be revolutionized by the sub-prime mortgage movement, merely a feeling that there was money to be made from backing new buyers. Hobart and Huber call these “mean-reverting bubbles,” presumably because there’s no expectation that the underlying developments would move the world forward. Fads merely rise and fall.
- On the other hand, Hobart and Huber call bubbles based on technological progress – as in the case of the railroads and the internet – “inflection bubbles.” After an inflection-driven bubble, the world will not revert to its prior state. In such a bubble, “investors decide that the future will be meaningfully different from the past and trade accordingly.” As Thompson tells us:

The definitive book on bubbles has long been Carlota Perez’s *Technological Revolutions and Financial Capital*. Bubbles were – are – thought to be something negative and to be avoided, particularly at the time Perez published her book. The year was 2002 and much of the world was in a recession coming off the puncturing of the dot-com bubble.

Perez didn’t deny the pain: in fact, she noted that similar crashes marked previous revolutions, including the Industrial Revolution, railways, electricity, and the automobile. In each case the bubbles were not regrettable, but necessary: **the speculative mania enabled what Perez called the “Installation Phase,” where necessary but not necessarily financially wise investments laid the groundwork for the “Deployment Period.” What marked the shift to the deployment period was the popping of the bubble; what enabled the deployment period were the money-losing investments.** (All emphasis added)

This distinction is very meaningful for Hobart and Huber, and I agree. They say, “not all bubbles destroy wealth and value. Some can be understood as important catalysts for techno-scientific progress.”

But I would restate as follows: “Mean-reversion bubbles” – in which markets soar on the basis of some new financial miracle and then collapse – destroy wealth. On the other hand, “inflection bubbles” based on revolutionary developments accelerate technological progress and create the foundation for a more prosperous future, **and they destroy wealth**. The key is to not be one of the investors whose wealth is destroyed in the process of bringing on progress.

Hobart and Huber go on to describe in greater depth the process through which bubbles finance the building of the infrastructure required by the new technology and thus accelerate its adoption:

Most novel technology doesn't just appear *ex nihilo* [i.e., from nothing], entering the world fully formed and all at once. Rather, it builds on previous false starts, failures, iterations, and historical path dependencies. Bubbles create opportunities to deploy the capital necessary to fund and speed up such large-scale experimentation – which includes lots of trial and error done in parallel – thereby accelerating the rate of potentially disruptive technologies and breakthroughs.

By generating positive feedback cycles of enthusiasm and investment, bubbles can be net beneficial. Optimism can be a self-fulfilling prophecy. Speculation provides the massive financing needed to fund highly risky and exploratory projects; what appears in the short term to be excessive enthusiasm or just bad investing turns out to be essential for bootstrapping social and technological innovations . . . A bubble can be a collective delusion, but it can also be an expression of collective vision. That vision becomes a site of coordination for people and capital and for the parallelization of innovation. Instead of happening over time, bursts of progress happen simultaneously across different domains. And with mounting enthusiasm . . . comes increased risk tolerance and strong network effects. The fear of missing out, or FOMO, attracts even more participants, entrepreneurs, and speculators, further reinforcing this positive feedback loop. Like bubbles, FOMO tends to have a bad reputation, but it's sometimes a healthy instinct. After all, none of us wants to miss out on a once-in-a-lifetime chance to build the future.

In other words, bubbles based on technological progress are good because they excite investors into pouring in money – a good bit of which is thrown away – to carpet-bomb a new area of opportunity and thus jump-start its exploitation.

The key realization seems to be that if people remained patient, prudent, analytical, and value-insistent, novel technologies would take many years and perhaps decades to be built out. Instead, the hysteria of the bubble causes the process to be compressed into a very short period – with some of the money going into life-changing investment in the winners but a lot of it being incinerated.

A bubble has aspects that are both technological and financial, but the above citations are from the standpoint of people who crave technological progress and are perfectly happy to see investors lose money in its interest. “We,” on the other hand, would like to see technological progress but have no desire to throw away money to help bring it about.

Ben Thompson ends this discussion by saying, “This is why I’m excited to talk about new technologies, the prospect for which *I don’t know*.” I love the fact that he’s excited by future possibilities and at the same time admits that the shape of the future is unknown (in our world, we might say “very risky”).

Assessing the Current Landscape

Now let’s get down to what we used to call “brass tacks.” What do we know? First, I haven’t met anyone who doesn’t believe artificial intelligence has the potential to be one of the biggest technological developments of all time, reshaping both daily life and the global economy.

We also know that in recent years, economies and markets have become increasingly dependent on AI:

- AI is responsible for a very large portion of companies' total capital expenditures.
- Capital expenditures on AI capacity account for a large share of the growth in U.S. GDP.
- AI stocks have been the source of the vast majority of the gains of the S&P 500.

As a *Fortune* headline put it on October 7:

75% of gains, 80% of profits, 90% of capex – AI's grip on the S&P is total and Morgan Stanley's top analyst is 'very concerned'

Further, I think it's important to note that whereas the gains in AI-related stocks account for a disproportionate percentage of the total gains in all stocks, the excitement AI injects into the market must have added a lot to the appreciation of non-AI stocks as well.

AI-related stocks have shown astronomical performance, led by Nvidia, the leading developer of computer chips for AI. From its formation in 1993 and its initial public offering in 1999, when its estimated market value was \$626 million, Nvidia briefly became the world's first company worth \$5 trillion. That's appreciation of around 8,000x, or roughly 40% a year for 26+ years. No wonder imaginations have been fired.

What Are the Areas of Uncertainty?

I think it's fair to say that while we know AI will be a source of incredible change, most of us have no idea exactly what it will be able to do, how it will be applied commercially, or what the timing will be.

Who will be the winners, and what will they be worth? If a new technology is assumed to be a world changer, it's invariably assumed that the leading companies possessing that technology will be of great value. But how accurate will that assumption prove to be? As Warren Buffett pointed out in 1999, "[The automobile was] the most important invention, probably, of the first half of the 20th century. . . . If you had seen at the time of the first cars how this country would develop in connection with autos, you would have said, 'This is the place I must be.' But of the 2,000 companies, as of a few years ago, only three car companies survived. So autos had an enormous impact on America but the opposite direction on investors." (*Time*, January 23, 2012)

In AI, there are some very strong leaders at present, including some of the world's strongest and richest companies. But new technology is notoriously disruptive. Will today's leaders prevail or give way to upstarts? How much will the arms race cost, and who will win?

Similarly, what's a share in an upstart worth? Unlike front runners worth trillions, it's possible to invest in some would-be challengers at enterprise values in mere billions or even – might I say? – millions. On June 25, 2024, CNBC reported as follows:

A team founded by college dropouts has raised \$120 million from investors led by Primary Venture Partners to build a new AI chip to take on Nvidia. Etched CEO Gavin Uberti said the startup is betting that as AI develops, most of the technology's power-hungry computing requirements will be filled by customized, hard-wired chips called ASICs. "If transformers go away, we'll die," Uberti told CNBC. "But if they stick around, we're the biggest company of all time."

Even granting the possibility that Etched won't become the biggest company of all time, if success could give them a valuation just one-fifth of Nvidia's peak – a mere \$1 trillion – what probability of success would be required to justify an investment of \$120 million? Assuming for simplicity's sake that the investment was for a 100% ownership stake, all you need is a belief that achieving the trillion-dollar value has a probability of one-tenth of a percent for an expected return of over eight times your money. Who's to say Etched doesn't have that chance? And in that case, why would anyone not play? The foregoing is what I call "lottery-ticket thinking," in which the dream of an enormous payoff justifies – no, compels – participation in an endeavor with an overwhelming probability of failing.

There's nothing wrong with calculating expected values this way. Leading venture capitalists engage in it every day to great effect. But assumptions regarding the possible payoffs and their probabilities must be reasonable. Thinking about a trillion-dollar payout will override reasonableness in any calculation.

Will AI produce profits, and for whom? Two things we know little or nothing about are the profits AI will produce for vendors and its impact on non-AI companies, primarily meaning those who employ it.

Will AI be a monopoly or duopoly, in which one or two leading companies are able to charge dearly for the capabilities? Or will it be a highly competitive free-for-all in which a number of firms compete on price for users' spending on AI services, making it a commodity? Or, perhaps most likely, will it be a mix of leading companies and specialized players, some of whom compete on price and others through proprietary advantages. It's said that the services currently responding to AI queries, such as ChatGPT and Gemini, lose money on every query they answer (of course, it's not unusual for participants in a new industry to offer "loss leaders" for a while). Will the leading tech firms – used to success in winner-take-all markets – be content to experience losses in their AI businesses for years in order to gain share? Hundreds of billions of dollars are being committed to the race for AI leadership. Who will win, and what will be the result?

Likewise, what will be AI's impact on the companies that use it? Clearly, AI will be a great tool for enhancing users' productivity by, among other things, replacing workers with computer-sourced labor and intelligence. But will this ability to cut costs add to the profit margins of the companies that employ it? Or will it simply enable price wars among those companies in the pursuit of customers? In that case, the savings might be passed on to the customers rather than garnered by the companies. In other words, is it possible AI will increase the efficiency of businesses without increasing their profitability?

Should we worry about so-called "circular deals"? In the telecom boom of the late 1990s, in which optical fiber became overbuilt, fiber-owning companies engaged in transactions with each other that permitted them to report profits. If two companies own fiber, they just have an asset on their books. But if each buys capacity from the other, they can both report profits . . . so they did. In other cases, manufacturers loaned network operators money to buy equipment from them, before the operators had customers to justify the buildout. All this resulted in profits that were illusory.

Nowadays, deals are being announced in which money appears to be round-tripped between AI players. People who believe there's an AI bubble find it easy to view these transactions with suspicion. Is the purpose to achieve legitimate business goals or to exaggerate progress?

Adding to worries, critics say, some of the deals that OpenAI has made with chipmakers, cloud computing companies and others are oddly circular. OpenAI is set to receive billions from tech companies but also sends billions back to the same companies to pay for computing power and other services. . . .

Nvidia has also made some deals that have raised questions about whether the company is paying itself. It announced that it would invest \$100 billion in OpenAI. The start-up receives that money as it buys or leases Nvidia's chips. . . .

Goldman Sachs has estimated that Nvidia will make 15 percent of its sales next year from what critics also call circular deals. (*The New York Times*, November 20)

Noteworthy, OpenAI has made investment commitments to industry counterparties totaling \$1.4 trillion, even though it has yet to turn a profit. The company makes clear that the investments are to be paid out of revenues received from the same parties and that it has ways to back out of these commitments. But all this raises the question of whether the AI industry has developed a perpetual motion machine.

(On this subject, I've been enjoying articles questioning the ability of people to relate to the word "trillion," and I think this idea is spot on. A million dollars is a dollar a second for 11.6 days. A billion dollars is a dollar a second for 31.7 years. We get that. But a trillion dollars is a dollar a second for 31,700 years. Who can get their head around the significance of 31,700 years?)

What will be the useful life of AI assets? We have to wonder whether the topic of obsolescence is being handled correctly in AI-land. What will be the lifespan of AI chips? How many years of earnings growth should be counted on in assigning p/e ratios for AI-related stocks? Will chips and other aspects of AI infrastructure last long enough to repay the debt undertaken to buy them? Will artificial general intelligence (a machine capable of doing anything the human brain can do) be achieved? Will that be the end of progress, or might there be further revolutions, and what firms will win them? Will firms reach a position where technology is stable and they can extract economic value from it? Or will new technologies continually threaten to supplant older ones as the route to success?

In this connection, a single issue of an FT newsletter briefly mentioned two developments that suggest the fluid nature of the competitive landscape:

- A study by the Massachusetts Institute of Technology and open-source AI start-up Hugging Face found that the total share of downloads of new Chinese-made open models rose to 17 per cent in the past year. The figure surpasses the 15.8 per cent share of downloads from American developers such as Google, Meta and OpenAI – the first time Chinese groups have beaten their American counterparts. . . .
- Nvidia shares fell sharply yesterday on fears that Google is gaining ground in artificial intelligence, erasing \$115bn in market value from the AI chipmaker. (*FirstFT Americas*, November 26)

Dynamic change creates the opportunity for incredible new technologies, but that same dynamism can threaten the leading companies' reign. Amid all these uncertainties, investors must ask whether the assumption of continued success incorporated in the prices they're paying is fully warranted.

Is exuberance leading to speculative behavior? For an extreme example, I'll cite the trend toward venture capital investments in startups via \$1 billion "seed rounds." Here's one vignette:

Thinking Machines, an AI startup helmed by former Open AI executive Mira Murati, just raised the largest seed round in history: \$2 billion in funding at a \$10 billion valuation. The company has not released a product and has refused to tell investors what they're even trying to build. "It was the most absurd pitch meeting," one investor who met with Murati said. "She was like, 'So we're doing an AI company with the best AI people, but

we can't answer any questions.' ” (“This Is How the AI Bubble Will Pop,” Derek Thompson Substack, October 2)

But that's ancient history. . . already two months old. Here's an update:

Thinking Machines Lab, the artificial intelligence startup founded by former Open AI executive Mira Murati, is in early talks to raise a new funding round at a roughly \$50 billion valuation, *Bloomberg News* reported on Thursday. The startup was last valued at \$12 billion in July, after it raised about \$2 billion. (*Reuters*, November 13)

And Thinking Machines Lab isn't alone:

In one of the boldest bets yet in the AI arms race, Safe Superintelligence (SSI), the stealth startup founded by former OpenAI chief scientist Ilya Sutskever, has raised \$2 billion in a round that values the company at \$32 billion – despite having no publicly released product or service. (*CTech by Calcalist*, April 13)

What's the end state? Part of the issue with AI includes the unusual nature of this newest thing. This isn't like a business that designs and sells a product, making money if the selling price exceeds the cost of the inputs. Rather, it's companies building an airplane while it's in flight, and once it's built, they'll know what it can do and whether anyone will pay for its services.

Many companies justify their spending because **they're not just building a product, they're creating something that will change the world:** artificial general intelligence, or A.G.I. . . . The rub is that none of them quite know how to do it.

But Anton Korinek, an economist at the University of Virginia, said the spending would all be justified if Silicon Valley reached its goal. He is optimistic it can be done.

“It's a bet on A.G.I. or bust,” Dr. Korinek said. (*The New York Times*, November 20 – emphasis added)

The yet-to-be-determined nature of the industry under construction is best captured in remarks from Sam Altman, the CEO of OpenAI, that have been paraphrased as follows: “we'll build this sort of generally intelligent system and then ask it to figure out a way to generate an investment return from it.”

This should be a source of pause for people who heretofore fully comprehended the nature of the businesses they invested in. Clearly, the value of a technology that equals or surpasses the human brain should be pretty big, but isn't it well beyond calculation?

A Word About the Use of Debt

To date, much of the investment in AI and the supporting infrastructure has consisted of equity capital derived from operating cash flow. But now, companies are committing amounts that require debt financing, and for some of those companies, the investments and leverage have to be described as aggressive.

The AI data centre boom was never going to be financed with cash alone. The project is too big to be paid for out of pocket. JPMorgan analysts have done some sums on the back of a napkin, or possibly a tablecloth, and estimated the bill for the infrastructure build-out

would come to \$5tn (not including a tip). Who knows if that's right, but we have good reason to expect close to half a trillion in spending next year. Meanwhile, the biggest spenders (Microsoft, Alphabet, Amazon, Meta and Oracle) had only about \$350bn in the bank, collectively, as of the end of the third quarter. ("Unhedged," *Financial Times*, November 13)

The firms mentioned above derive healthy cash flows from their very strong non-AI businesses. But the massive, winner-take-all arms race in AI is requiring some to take on debt. **In fact, it's reasonable to think one of the reasons they're spending vast sums is to make it hard for lesser firms to keep up.**

Oracle, Meta, and Alphabet have issued 30-year bonds to finance AI investments. In the case of the latter two, the yields on the bonds exceed those on Treasuries of like maturity by 100 basis points or less. Is it prudent to accept 30 years of technological uncertainty to make a fixed-income investment that yields little more than riskless debt? And will the investments funded with debt – in chips and data centers – maintain their level of productivity long enough for these 30-year obligations to be repaid?

On November 14, Alex Kantrowitz's *Big Technology Podcast* carried a conversation with Gil Luria, Head of Technology Research at financial services firm D.A. Davidson, primarily regarding the use of debt in the AI sector. Here's some of what Luria had to say:

- Healthy behavior is being practiced by "... reasonable, thoughtful business leaders, like the ones at Microsoft, Amazon, and Google that are making sound investments in growing the capacity to deliver AI. And the reason they can make sound investments is that they have all the customers. ... And so, when they make investments, they're using cash on their balance sheets; they have tremendous cash flow to back it up; they understand that it's a risky investment; and they balance it out."
- Unhealthy behavior – Here he describes "... a startup that is borrowing money to build data centers for another startup. They're both losing tremendous amounts of cash, and yet they're somehow being able to raise this debt capital in order to fund this buildout, again without having the customers or the visibility into those investments paying off."
- "So there's a whole range of behaviors between healthy and unhealthy, and we just need to sort that out so we don't make the mistakes of the past."
- "There are certain things we finance through equity, through ownership, and there are certain things we finance through debt, through an obligation to pay down interest over time. And as a society, for the longest time, we've had those two pieces in their right place. Debt is when I have a predictable cash flow and/or an asset that can back that loan, and then it makes sense for me to exchange capital now for future cash flows to the lender. ... We use equity for investing in more speculative things, for when we want to grow and we want to own that growth, but we're not sure about what the cash flow is going to be. That's how a normal economy functions. When you start confusing the two you get yourself in trouble."

Among potentially worrisome factors, Luria cites these:

- "A speculative asset ... we don't know how much of it we're really going to need in two to five years."
- Lender personnel with incentives to make loans but no exposure to long-term consequences
- The possibility that the supply of AI capacity catches up with or surpasses the demand
- The chance that future generations of AI chips will be more powerful, obsoleting existing ones or reducing their value as backing for debt
- Powerful competitors who vie for market share by cutting rental rates and running losses

Here are some important paragraphs from Azeem Azhar's *Exponential View* of October 18:

When does an AI boom tip into a bubble? [Investor and engineer] Paul Kedrosky points to the Minsky moment – the inflection point when credit expansion exhausts its good projects and starts chasing bad ones, funding marginal deals with vendor financing and questionable coverage ratios. For AI infrastructure, that shift may already be underway; the telltale signs include hyperscalers' capex outpacing revenue momentum and lenders sweetening terms to keep the party alive.

Paul makes a compelling case. We've entered speculative finance territory – arguably past the tentative stage – and recent deals will set dangerous precedents. As Paul warns, this financing will “create templates for future such transactions,” spurring rapid expansion in junk issuance and SPV proliferation among hyperscalers chasing dominance at any cost. . . .

For AI infrastructure, the warning signs are flashing: vendor financing proliferates, coverage ratios thin, and hyperscalers leverage balance sheets to maintain capex velocity even as revenue momentum lags. We see both sides – genuine infrastructure expansion alongside financing gymnastics that recall the 2000 telecom bust. The boom may yet prove productive, but only if revenue catches up before credit tightens. When does healthy strain become systemic risk? That's the question we must answer before the market does. (Emphasis added)

Azhar references the use of off-balance sheet financing via special-purpose vehicles, or SPVs, which were among the biggest contributors to Enron's precariousness and eventual collapse. A company and its partners set up an SPV for some specific purpose(s) and supply the equity capital. The parent company may have operating control, but because it doesn't have majority ownership, it doesn't consolidate the SPV on its financial statements. The SPV takes on debt, but that debt doesn't appear on the parent's books. The parent may be an investment grade borrower, but likewise, the debt isn't an obligation of the parent or guaranteed by it. Today's debt may be backed by promised rent from a data center tenant – sometimes an equity partner – but the debt isn't a direct obligation of the equity partner either. Essentially, an SPV is a way to make it look like a company isn't doing the things the SPV is doing and doesn't have the debt the SPV does. (Private equity funds and private credit funds are highly likely to be found among the partners and lenders in these entities.)

As I quoted earlier, according to Perez (who wrote on the heels of the dot-com bubble), “what enabled the deployment period were the money-losing investments.” Early investment is lost in the “Minsky moment,” in which unwise commitments made in an extended up-cycle encounters value destruction in a correction. And there are three things we know for sure about the use of debt:

- it magnifies losses if there are losses (just as it magnifies the hoped-for gains if they materialize),
- it increases the probability of a venture failing if it encounters a difficult moment, and
- despite the layer of equity beneath it, it puts lenders' capital at risk if the difficult moment is bad enough.

One key risk to consider is the possibility that the boom in data center construction will result in a glut. Some data centers may be rendered uneconomic, and some owners may go bankrupt. In that case, a new generation of owners might buy up centers at pennies on the dollar from lenders who foreclosed on them, reaping profits when the industry stabilizes. This is a process through which “creative destruction” brings markets into equilibrium and reduces costs to levels that make future business profitable.

Debt is neither a good thing nor a bad thing *per se*. Likewise, the use of leverage in the AI industry shouldn't be applauded or feared. It all comes down to the proportion of debt in the capital structure; the quality of the assets or cash flows you're lending against; the borrowers' alternative sources of liquidity for repayment; and the adequacy of the safety margin obtained by lenders. We'll see which lenders maintain discipline in today's heady environment.

It's worth noting in this connection that Oaktree has made a few investments in data centers, and our parent, Brookfield, is raising a \$10 billion fund for investment in AI infrastructure. Brookfield is putting up its own money and has equity commitments from sovereign wealth funds and Nvidia, to which it intends to apply "prudent" debt. Brookfield's investments seem likely to go largely into geographies that are less saturated with data centers and for infrastructure to supply the vast amounts of electric power that data centers will require. Of course, we're both doing these things on the basis of what we think are prudent decisions.

I know I don't know enough to opine on AI. But I do know something about debt, and it's this:

- **It's okay to supply debt financing for a venture where the outcome is uncertain.**
- **It's not okay where the outcome is purely a matter of conjecture.**
- **Those who understand the difference still have to make the distinction correctly.**

The FT's *Unhedged* quotes Chong Sin, lead analyst for CMBS research at JPMorgan, as saying, "... in our conversations with investment grade ABS and CMBS investors, one often-cited concern is whether they want to take on the residual value risk of data centers when the bonds mature." I'm glad potential lenders are asking the kind of questions they should.

Here's how to think about the intersection of debt and AI according to Bob O'Leary, Oaktree's co-CEO and co-portfolio manager of our Opportunities Funds:

Most technological advances develop into winner-takes-all or winner-takes-most competitions. The "right" way to play this dynamic is through equity, not debt.

Assuming you can diversify your equity exposures so as to include the eventual winner, the massive gain from the winner will more than compensate for the capital impairment on the losers. That's the venture capitalist's time-honored formula for success.

The precise opposite is true of a diversified pool of debt exposures. You'll only make your coupon on the winner, and that will be grossly insufficient to compensate for the impairments you'll experience on the debt of the losers.

Of course, if you can't identify the pool of companies from which the winner will emerge, the difference between debt and equity is irrelevant – you're a zero either way. I mention this because that's precisely what happened in search and social media: early leaders (Lycos in search and MySpace in social media) lost out spectacularly to companies that emerged later (Google in search and Facebook in social media).

Trying to Get to a Conclusion

There can be no doubt that today's behavior is "speculative," defined as based on speculation regarding the future. There's also no doubt that no one knows what the future holds, but investors are betting huge sums on that future.

In that connection, I want to say a little about the unique nature of AI. The AI revolution is different from the technological revolutions that preceded it in ways that are both wonderful and worrisome. It feels to me like a genie has been released from a bottle, and it isn't going back in:

AI may not be a tool for mankind, but rather something of a replacement. It may be capable of taking over cognition, on which humans have thus far had a monopoly. Because of this, it's likely to be different in kind from prior developments, not just in degree. (More on this in my postscript.)

AI technology is progressing at an incredibly rapid clip, possibly leaving scant time for mankind to adjust. I'll provide two examples:

- Coding, which we called “computer programming” 60 years ago, is the canary in the coal mine in terms of the impact of AI. In many advanced software teams, developers no longer write the code; they type in what they want, and AI systems generate the code for them. Coding performed by AI is at a world-class level, something that wasn't so just a year ago. According to my guide here, “There is no speculation about whether or not human replacement will take place in that vertical.”
- In the field of digital advertising, when users log into an app, AI engages in “ad matching,” showing them ads tailored to the preferences displayed by their prior surfing. No humans need apply to do this job.

Perhaps most importantly, **the growth of demand for AI seems totally unpredictable**. As one of my younger advisers explained, “the speed and scale of improvement mean it's incredibly hard to forecast demand for AI. Adoption today may have nothing to do with adoption tomorrow, because a year or two from now, AI may be able to do 10x or 100x what it can do today. Thus, how can anyone say how many data centers will be needed? And how can even successful companies know how much computing capacity to contract for?”

With differences like these, how can anyone correctly judge what AI implies for the future?

* * *

One of the things occupying many observers at this juncture – including me – is the search for parallels to past bubbles. Here's some historical perspective from a recent article in *Wired*:

AI's closest historical analogue here may be not electric lighting but radio. When RCA started broadcasting in 1919, it was immediately clear that it had a powerful information technology on its hands. But less clear was how that would translate into business. “Would radio be a loss-leading marketing for department stores? A public service for broadcasting Sunday sermons? An ad-supported medium for entertainment?” [Brent Goldfarb and David A. Kirsch of the University of Maryland] write. “All were possible. All were subjects of technological narratives.” As a result, radio turned into one of the biggest bubbles in history – peaking in 1929, before losing 97 percent of its value in the crash. This wasn't an incidental sector; RCA was, along with Ford Motor Company, the most high-traded stock on the market. It was, as *The New Yorker* recently wrote, “the Nvidia of its day.” . . .

In 1927, Charles Lindbergh flew the first solo nonstop transatlantic flight from New York to Paris. . . . It was the biggest tech demo of the day, and it became an enormous,

ChatGPT-launch-level coordinating event – a signal to investors to pour money into the industry.

“Expert investors appreciated correctly the importance of airplanes and air travel,” Goldfarb and Kirsch write, but **“the narrative of inevitability largely drowned out their caution. Technological uncertainty was framed as opportunity, not risk. The market overestimated how quickly the industry would achieve technological viability and profitability.”**

As a result, the bubble burst in 1929 – from its peak in May, aviation stocks dropped 96 percent by May 1932. . . .

It’s worth reiterating that two of the closest analogs AI seems to have in tech bubble history are aviation and broadcast radio. Both were wrapped in high degrees of uncertainty and both were hyped with incredibly powerful coordinating narratives. Both were seized on by pure play companies seeking to capitalize on the new game-changing tech, and both were accessible to the retail investors of the day. Both helped inflate a bubble so big that when it burst, in 1929, it left us with the Great Depression. (“AI Is the Bubble to Burst Them All,” Brian Merchant, *Wired*, October 27 – emphasis added. N.b., the Depression had many causes beyond the bursting of the radio/aviation bubble.)

Derek Thompson, who supplied the quote with which I opened this memo, ended his newsletter with some terrific historical perspective:

The railroads were a bubble and they transformed America. Electricity was a bubble, and it transformed America. The broadband build-out of the late-1990s was a bubble that transformed America. I am not rooting for a bubble, and quite the contrary, I hope that the US economy doesn’t experience another recession for many years. **But given the amount of debt now flowing into AI data center construction, I think it’s unlikely that AI will be the first transformative technology that isn’t overbuilt and doesn’t incur a brief painful correction.** (“AI Could Be the Railroad of the 21st Century. Brace Yourself.” November 4 – emphasis added)

The skeptics readily cite ways in which today’s events are comparable to the internet bubble:

- A change-the-world technology
- Exuberant, speculative behavior
- The role of FOMO
- Suspect, circular deals
- The use of SPVs
- \$1 billion seed rounds

The supporters have reasons why the comparison isn’t appropriate:

- An existing product for which there is strong demand
- One billion users already (many times the number of internet users at the height of the bubble)
- Well-established main players with revenues, profits, and cash flow
- The absence of an IPO craze with prices doubling in a day
- Reasonable p/e ratios for the established participants

I'll elaborate regarding the first of the proposed non-comparable factors. Unlike in the internet bubble, AI products already exist at scale, the demand for them is exploding, and they're producing revenues in rapidly increasing amounts. For example, Anthropic, one of the two leaders in producing models for AI coding as described on page 12, is said to have "10x-ed" its revenues in each of the last two years (for those who didn't study higher math, that's 100x in two years). Revenues from Claude Code, a program for coding that Anthropic introduced earlier this year, already are said to be running at an annual rate of \$1 billion. Revenues for the other leader, Cursor, were \$1 million in 2023 and \$100 million in 2024, and they, too, are expected to reach \$1 billion this year.

As to the final bullet point, see the table below, which comes from Goldman Sachs via Derek Thompson. You'll notice that during the internet bubble of 1998-2000, the p/e ratios were much higher for Microsoft, Cisco, and Oracle than they are today for the biggest AI players – Nvidia, Microsoft, Alphabet, Amazon, and Meta (OpenAI doesn't have earnings). In fact, Microsoft's on a half-off sale relative to its p/e 26 years ago! In the first bubble I witnessed – surrounding the Nifty-Fifty in 1969-72 – the p/e ratios for the leading companies were even higher than those of 1998-2000.

Exhibit 7: Valuation of largest TMT stocks today vs. Tech Bubble

Dec-99			Current		
Ticker	Wgt in SPX	NTM P/E	Ticker	Wgt in SPX	NTM P/E
Microsoft Corp.	4.9%	69 x	NVIDIA Corp.	7.6%	30 x
Cisco Systems Inc.	2.8	101	Microsoft Corp.	6.9	32
Intel Corp.	2.2	32	Apple Inc.	6.5	30
Lucent Tech. Inc.	1.9	49	Alphabet Inc.	2.5	22
IBM Corp.	1.6	28	Amazon.com, Inc.	3.9	31
Time Warner Inc.	1.4	217	Meta Platforms Inc.	2.9	25
AT&T Inc.	1.4	22	Broadcom Inc.	2.6	38
AT&T Corp.	1.3	24	Tesla, Inc.	1.7	151
Oracle Corp.	1.3	90	Visa Inc.	1.1	28
WorldCom Inc.	1.2	30	Netflix, Inc.	1.0	40
Median		41 x	Median		31 x

Source: Goldman Sachs Global Investment Research

In Conclusion

For my final citation, I'll look to Sam Altman of OpenAI. His comments seem to me to capture the essence of what's going on:

"When bubbles happen, smart people get overexcited about a kernel of truth," Mr. Altman told reporters this year. "Are we in a phase where investors as a whole are overexcited about A.I.? My opinion is yes. Is A.I. the most important thing to happen in a very long time? My opinion is also yes." (*The New York Times*, November 20)

But do I have a bottom line? Yes, I do. Alan Greenspan's phrase, mentioned earlier, serves as an excellent way to sum up a stock market bubble: "irrational exuberance." **There is no doubt that investors are applying exuberance with regard to AI. The question is whether it's irrational. Given the vast potential of AI but also the large number of enormous unknowns, I think virtually no one can say for sure.** We can theorize about whether the current enthusiasm is excessive, but we won't know until years from now whether it was. Bubbles are best identified in retrospect.

While the parallels to past bubbles are inescapable, believers in the technology will argue that “this time it’s different.” Those four words are heard in virtually every bubble, explaining why the present situation isn’t a bubble, unlike the analogous prior ones. On the other hand, Sir John Templeton, who in 1987 drew my attention to those four words, was quick to point out that 20% of the time things really are different. **But on the third hand, it must be borne in mind that behavior based on the belief that it’s different is what causes it to not be different!**

Today’s situation calls to mind a comment attributed to American economist Stuart Chase about faith. I believe it’s also applicable to AI (as well as to gold and cryptocurrencies):

For those who believe, no proof is necessary. For those who don't believe, no proof is possible.

Here’s my actual bottom line:

- There’s a consistent history of transformational technologies generating excessive enthusiasm and investment, resulting in more infrastructure than is needed and asset prices that prove to have been too high. The excesses accelerate the adoption of the technology in a way that wouldn’t occur in their absence. The common word for these excesses is “bubbles.”
- AI has the potential to be one of the greatest transformational technologies of all time.
- As I wrote just above, AI is currently the subject of great enthusiasm. If that enthusiasm doesn’t produce a bubble conforming to the historical pattern, that will be a first.
- Bubbles created in this process usually end in losses for those who fuel them.
- The losses stem largely from the fact that the technology’s newness renders the extent and timing of its impact unpredictable. This in turn makes it easy to judge companies too positively amid all the enthusiasm and difficult to know which will emerge as winners when the dust settles.
- There can be no way to participate fully in the potential benefits from the new technology without being exposed to the losses that will arise if the enthusiasm and thus investors’ behavior prove to have been excessive.
- The use of debt in this process – which the high level of uncertainty usually precluded in past technological revolutions – has the potential to magnify all of the above this time.

Since no one can say definitively whether this is a bubble, I’d advise that no one should go all-in without acknowledging that they face the risk of ruin if things go badly. But by the same token, no one should stay all-out and risk missing out on one of the great technological steps forward. A moderate position, applied with selectivity and prudence, seems like the best approach.

Finally, it’s essential to bear in mind that there are no magic words in investing. These days, people promoting real estate funds say, “Office buildings are so yesterday, but we’re investing in the future through data centers,” whereupon everyone nods in agreement. But data centers can be in shortage or in oversupply, and rental rates can surprise to the upside or the downside. As a result, they can be profitable . . . or not. **Intelligent investment in data centers, and thus in AI – like everything else – requires sober, insightful judgment and skillful implementation.**

December 9, 2025

P.S.: The following has nothing to do with the financial markets or the question of whether AI is the subject of a bubble. My topic is the impact of AI on society through joblessness and purposelessness.

You needn't read it – that's why it's a postscript – but it's important to me, and I've been looking for a place to say a few words about it.

On November 18, a research note from Barclays described Fed Governor Christopher Waller as having “highlighted how recent stock market enthusiasm around AI has not yet translated into job creation.” This strikes me as paradoxical given my sense that one of AI's main impacts will be to increase productivity and thus eliminate jobs. That is the source of my concern.

I view AI primarily as an incredible labor-saving device. Joe Davis, Global Chief Economist and Global Head of the Investment Strategy Group at Vanguard, says, “for most jobs – likely four out of five – AI's impact will result in a mixture of innovation and automation, and could save about 43% of the time people currently spend on their work tasks.” (*Exponential View*, September 3)

I find the resulting outlook for employment terrifying. I am enormously concerned about what will happen to the people whose jobs AI renders unnecessary, or who can't find jobs because of it. The optimists argue that “new jobs have always materialized after past technological advances.” I hope that'll hold true in the case of AI, but hope isn't much to hang one's hat on, and I have trouble figuring out where those jobs will come from. **Of course, I'm not much of a futurist or a financial optimist, and that's why it's a good thing I shifted from equities to bonds in 1978.**

The other thing the optimists say is that “the beneficial impact of AI on productivity will cause a huge acceleration in GDP growth.” Here I have specific quibbles:

- The change in GDP can be thought of as the change in hours worked times the change in output per hour (aka “productivity”). The role of AI in increasing productivity means it will take fewer hours worked – meaning fewer workers – to produce the goods we need.
- Or, viewed from the other direction, maybe the boom in productivity will mean a lot more goods can be produced with the same amount of labor. But if a lot of jobs are lost to AI, how will people be able to afford the additional goods AI enables to be produced?

I find it hard to imagine a world in which AI works shoulder-to-shoulder with all the people who are employed today. How can employment not decline? AI is likely to replace large numbers of entry-level workers, people who process paper without applying judgment, and junior lawyers who scour the lawbooks for precedents. Maybe even junior investment analysts who create spreadsheets and compile presentation materials. It's said that AI can read an MRI better than the average doctor. Driving is one of the most populous professions in America, and driverless vehicles are already arriving; where will all the people who currently drive taxis, limos, buses, and trucks find jobs?

I imagine government's response will be something called “universal basic income.” The government will simply mail checks to the millions for whom there are no jobs. But the worrier in me finds problems in this, too:

- **Where will the money come from for those checks?** The job losses I foresee imply reduced income tax receipts and increased spending on entitlements. This puts a further burden on the declining segment of the population that is working and implies even greater deficits ahead. In this new world, will governments be able to fund ever-increasing deficits?
- And more importantly, **people get a lot more from jobs than just a paycheck.** A job gives them a reason to get up in the morning, imparts structure to their day, gives them a productive role in society and self-respect, and presents them with challenges, the overcoming of which provides satisfaction. How will these things be replaced? I worry about large numbers of people receiving

subsistence checks and sitting around idle all day. I worry about the correlation between the loss of jobs in mining and manufacturing in recent decades and the incidence of opioid addiction and shortening of lifespans.

And by the way, if we eliminate large numbers of junior lawyers, analysts, and doctors, where will we get the experienced veterans capable of solving serious problems requiring judgment and pattern recognition honed over decades?

What jobs won't be eliminated? What careers should our children and grandchildren prepare for? Think about the jobs that machines can't perform. My list starts with plumbers, electricians, and masseurs – physical tasks. Maybe nurses will earn more than doctors because they deliver hands-on care. And what distinguishes the best artists, athletes, doctors, lawyers, and hopefully investors? I think it's something called talent or insight, which AI might or might not be able to replicate. But how many people at the top of those professions are needed? A past presidential candidate said he would give laptops to everyone who lost their job to offshoring. How many laptop operators do we need?

Finally, I'm concerned that a small number of highly educated multi-billionaires living on the coasts will be viewed as having created technology that puts millions out of work. **This promises even more social and political division than we have now, making the world ripe for populist demagoguery.**

I've seen incredible progress over the course of my lifetime, but in many ways I miss the simpler world I grew up in. I worry that this will be another big one. **I get no pleasure from this recitation. Will the optimists please explain why I'm wrong?**

Interestingly in this connection, Vanguard's Joe Davis points out that more Americans are turning 65 in 2025 than in any preceding year, and that approximately 16 million baby boomers will retire between now and 2035. Could AI merely make up for that? There's an optimistic take for you.

HM

Legal Information and Disclosures

This memorandum expresses the views of the author as of the date indicated and such views are subject to change without notice. Oaktree has no duty or obligation to update the information contained herein. Further, Oaktree makes no representation, and it should not be assumed, that past investment performance is an indication of future results. Moreover, wherever there is the potential for profit there is also the possibility of loss.

This memorandum is being made available for educational purposes only and should not be used for any other purpose. The information contained herein does not constitute and should not be construed as an offering of advisory services or an offer to sell or solicitation to buy any securities or related financial instruments in any jurisdiction. Certain information contained herein concerning economic trends and performance is based on or derived from information provided by independent third-party sources. Oaktree Capital Management, L.P. (“Oaktree”) believes that the sources from which such information has been obtained are reliable; however, it cannot guarantee the accuracy of such information and has not independently verified the accuracy or completeness of such information or the assumptions on which such information is based.

This memorandum, including the information contained herein, may not be copied, reproduced, republished, or posted in whole or in part, in any form without the prior written consent of Oaktree.

